

# **TACO2+ over Combat Net Radio (MIL-STD-188-220A)**

**John C. C. White**

**June 1996**

# Introduction

- 0 **There is a need for a bulk data transfer mechanism for the Combat Net Radio protocol suite**
- 0 **NETBLT has been identified as a candidate for this role**
- 0 **There has been considerable real-world experience with NETBLT, because it is the heart of TACO2, the tactical protocol component of the NITFS**
- 0 **This briefing describes:**
  - **the basic operation of the NETBLT protocol**
  - **the results of initial tests conducted using that protocol over CNR at the Digital Integration Lab, Ft. Monmouth**

# Some Elements of NETBLT

- 0 **NETBLT is a bulk data transmission protocol**
- 0 **Data is divided into “buffers”, which are further divided into (and transmitted as) packets**
  - **Buffers are acknowledged, packets are retransmitted**
  - **Multiple buffers may be outstanding concurrently**
- 0 **Values of several variables are initially negotiated**
- 0 **Operation is driven mainly from the receiver end**
  - **Reduces sensitivity to network delay**
  - **Retransmission timer is a critical element**
- 0 **Burst size and burst interval provide a throttling mechanism**
- 0 **There is a single control packet, which can hold multiple control messages**

# OPEN and RESPONSE Packet Format

## 5.2.6.1 OPEN (type 0) and RESPONSE (type 1):

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Checksum										Version										Type																			
Length										Local Port																													
Foreign Port										Longword Alignment Padding																													
Connection Unique ID																																							
Buffer Size																																							
DATA packet size										Burst Size																													
Burst Interval										Death Timer Value																													
Reserved (MBZ)										C	M	Maximum # Outstanding Buffers																											
Client String ...																																							
Longword Alignment Padding																																							

# DATA Packet Format

## 5.2.6.4 DATA (type 5) and LDATA (type 6):

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Checksum										Version										Type																			
Length										Local Port																													
Foreign Port										Longword Alignment Padding																													
Buffer Number																																							
Last Buffer Touched																																							
High Consecutive Seq Num Rcvd										Packet Number																													
Data Area Checksum Value										Reserved (MBZ)										L																			
New Burst Size										New Burst Interval																													

# CONTROL Packet Header Format

## 5.2.6.6 CONTROL (type 8):

0										1										2										3																																							
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																						
+-----+-----+-----+-----+																																																																					
										Checksum																				Version																				Type																			
+-----+-----+-----+-----+																																																																					
										Length																				Local Port																																							
+-----+-----+-----+-----+																																																																					
										Foreign Port																				Longword Alignment Padding																																							
+-----+-----+-----+-----+																																																																					

# GO and OK Message Format

## 5.2.6.6.1 GO message (type 0):

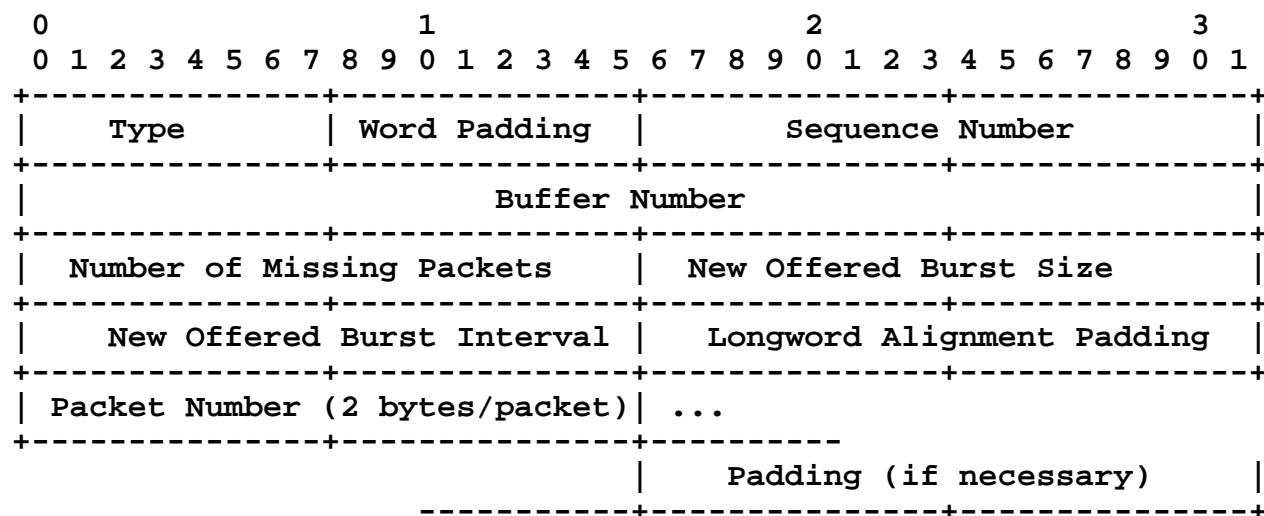
0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----+										+-----+										+-----+																			
Type										Word Padding										Sequence Number																			
+-----+										+-----+										+-----+																			
										Buffer Number																													
+-----+										+-----+										+-----+																			

## 5.2.6.6.2 OK message (type 1):

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----+										+-----+										+-----+																			
Type										Word Padding										Sequence Number																			
+-----+										+-----+										+-----+																			
										Buffer Number																													
+-----+										+-----+										+-----+																			
New Offered Burst Size										New Offered Burst Interval																													
+-----+										+-----+										+-----+																			
Current control timer value										Longword Alignment Padding																													
+-----+										+-----+										+-----+																			

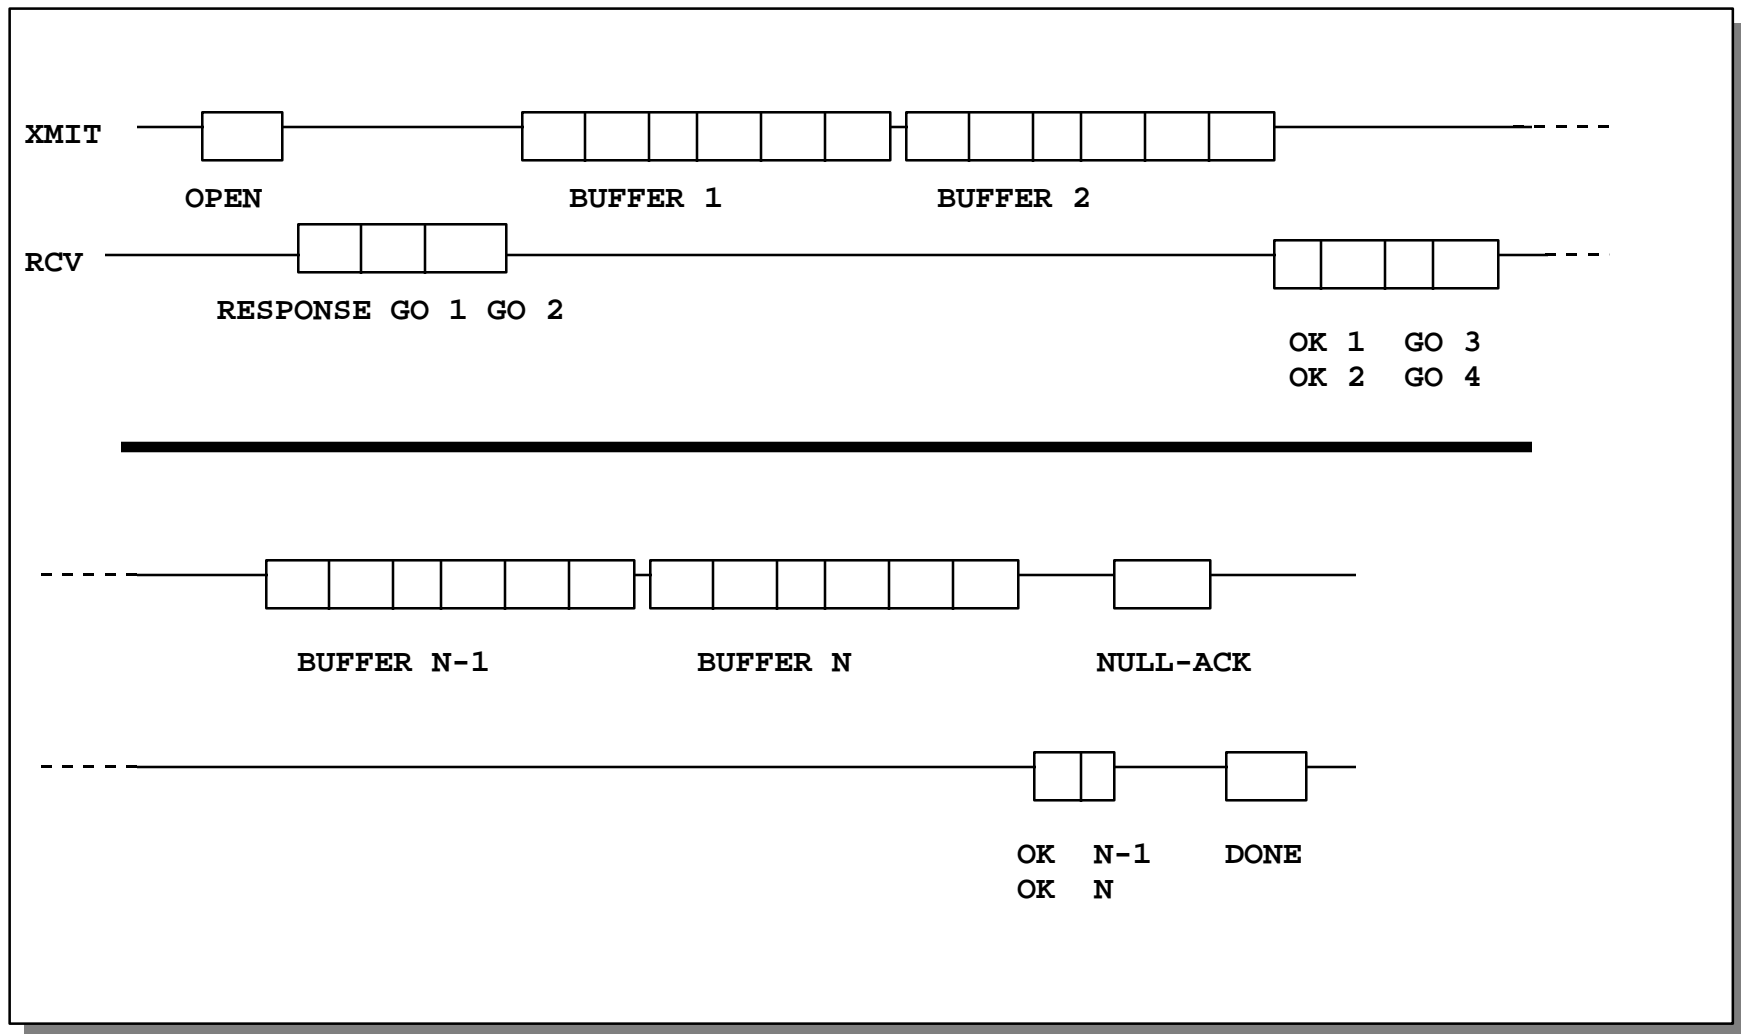
# RESEND Message Format

## 5.2.6.6.3 RESEND Message (type 2):






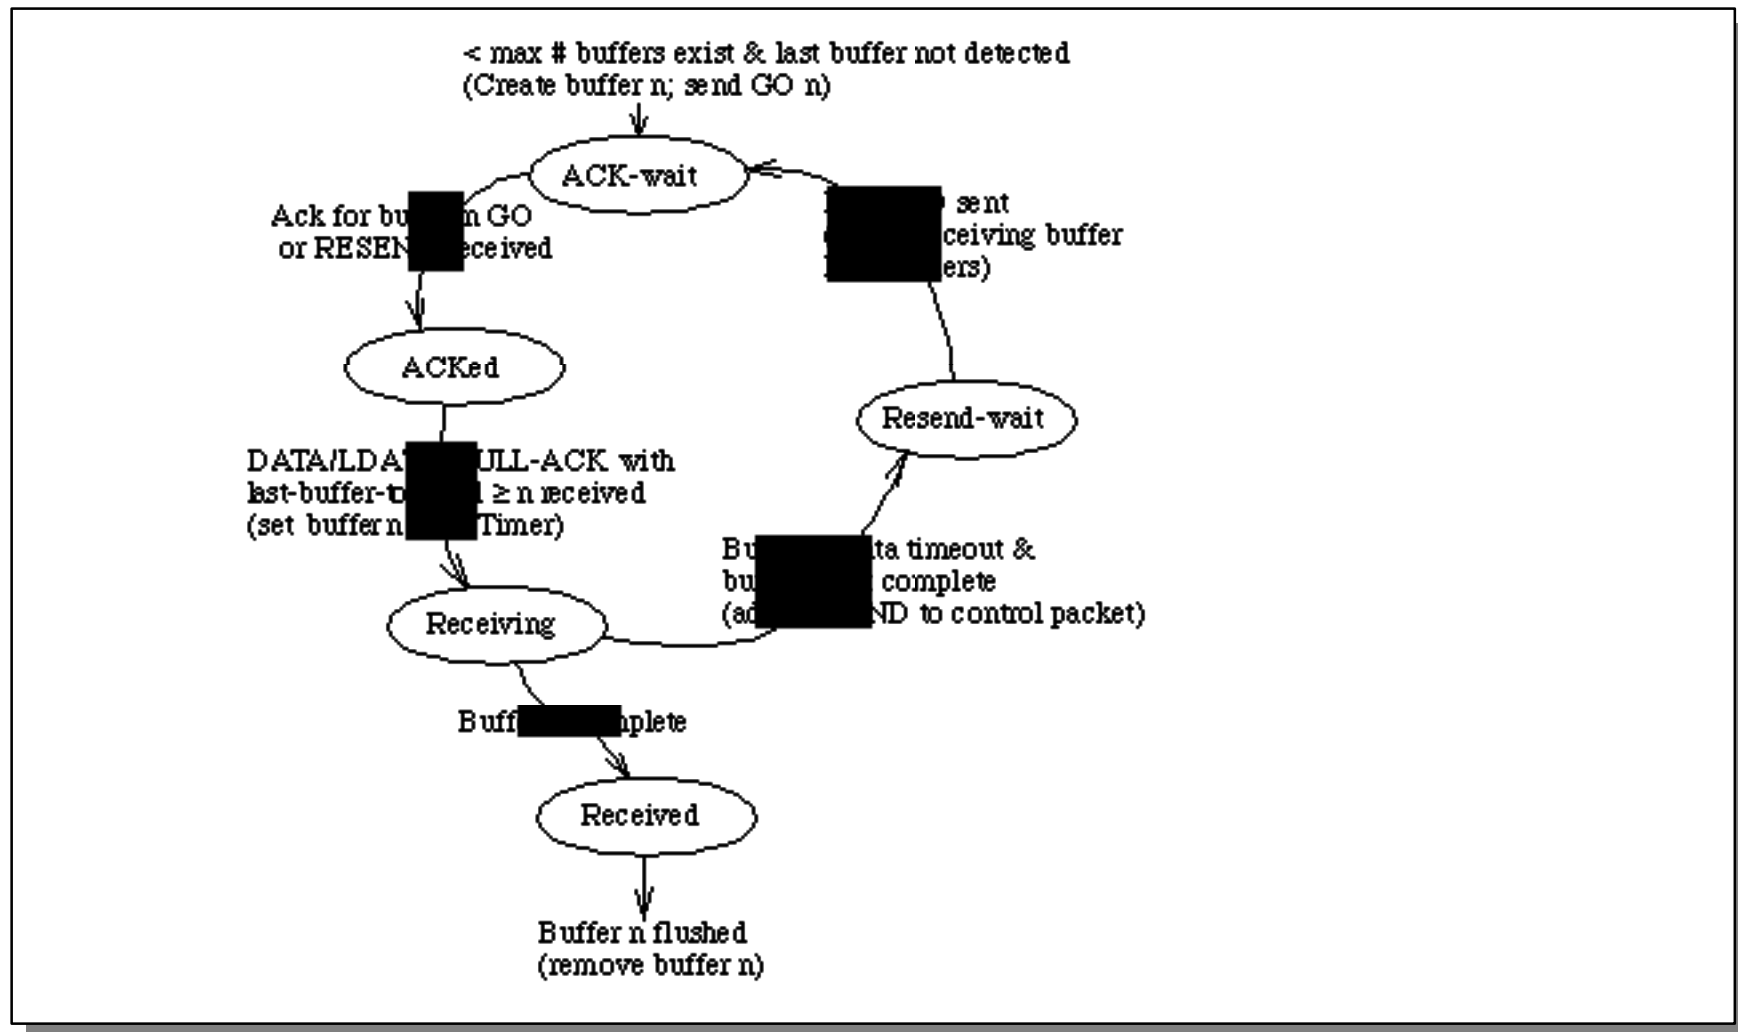
# NETBLT Operation in “Half-duplex” Mode



# Transmit Buffer State Diagram



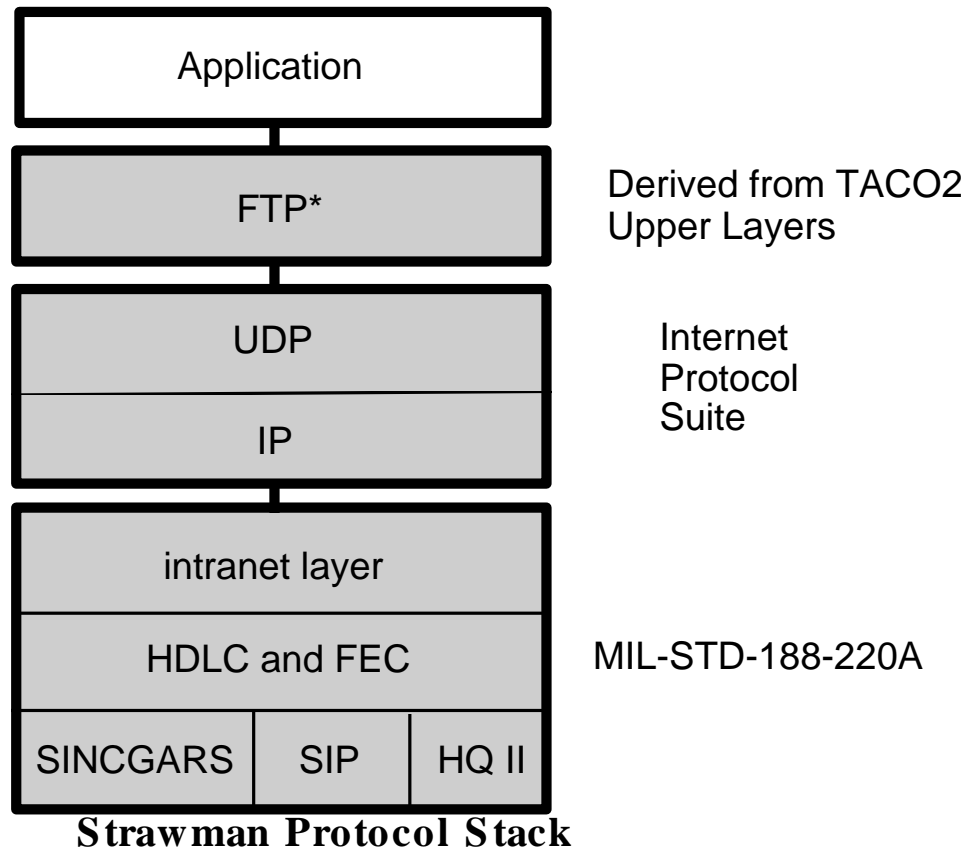
# Receive Buffer State Diagram




# Initial Testing

- 0 **Tests were conducted at the Digital Integration Lab at Ft. Monmouth NJ**
- 0 **Network had just two nodes, hard-wired connectivity**
- 0 **Normal MIL-STD-188-220A net access algorithms and FEC used, but no encryption**
- 0 **Equipment was newly delivered to the DIL**

# Protocol Stack Tested



# Test Hardware Configuration



# Initial Test Configuration

## 0 At each node:

- **MS-DOS host computer**
  - = **LSL.COM** - Novell Netware Link Support Layer v2.05
  - = **SLIP\_PPP.COM** - Novell Asynchronous SLIP\_PPP Driver v4.12
  - = **PPPPKT.COM** - Packet Driver interface to ODI v0.60
  - = **PACKET.EXE** - TACO2 Sample Software for MS-DOS v2.1beta
- **Internet Controller (INC)**
- **SINGCARS SIP radio**
- **radio handset**
- **appropriate interconnecting cables**

# Test Observations

- 0 **Optimum packet size = 480 bytes (420 data, 60 header)**
  - **Allows three packets per transmission, usually**
- 0 **Good buffer size = 4200 bytes, three buffers outstanding**
- 0 **Half-duplex operation of NETBLT preferred over full-duplex**
  - **Fewer receiver to sender transmissions**
  - **Fewer radio collisions**
- 0 **Presence or absence of UDP layer had negligible performance impact**
- 0 **Local data rate > 19200 caused some PPP checksum errors**
  - **Processor and software dependent**



## Test Observations (cont)

- 0 **The INC uses the IP TOS field to control the operation of the link layer. We were able to operate the link in the following modes:**
  - **a. Type 1 datagram, no link-layer acknowledgement**
  - **b. Type 1 datagram, immediate ("coupled") transmission of a link-layer acknowledgement for each packet**
  - **c. Type 4 datagram, transmission of a link-layer acknowledgement for each transmission burst**
- 0 **As expected, mode a was the most efficient in terms of throughput over a clean link, mode c was next most efficient, and mode b was the slowest, since it requires that the link be turned around for every packet**

## Test Observations (concluded)

- 0 **NETBLT Burst Interval set to fixed value worked best**
  - **INC delivers packets in bursts, which misled the current TACO2 packet timing algorithm**
  - **Improved packet timing algorithm will probably combine information from NETBLT Burst Interval and measured values with variance**
- 0 **Brief network disturbances (generated by keying microphone or passing traffic across the second INC ports) had little effect**
- 0 **Longer disturbances caused some retransmissions if packets were held up for too long**

# Test Setup

- 0 **Used three test files:**
  - **SMALL - 1000 bytes**
  - **MEDIUM - 10000 bytes**
  - **LARGE - 50000 bytes**
- 0 **The tests were run with these initial parameters:**
  - **NETBLT packet size = 420 bytes**
  - **NETBLT buffer size = 4200 bytes**
  - **# of NETBLT buffers = 3**
  - **Burst interval = 10 seconds**
  - **Burst size = 10 packets**
  - **PPP rate = 38400**
  - **IP TOS = 0x10 (Routine precedence, D bit set = low delay)**

# Test Results

0 File	Time (sec)	Rate (bps)
0 SMALL	7.52	1063
	7.63	1048
	8.73	916
0 MEDIUM	39.50	2025
	32.52	2460
	43.50	1839
	44.55	1796
0 In each of the last two runs, we observed that a control packet was lost and had to be retransmitted. We then reduced the PPP speed to 19200 bps and continued testing.		

## Test Results (cont)

0	<b>LARGE</b>	<b>160.52</b>	<b>2492</b>
		<b>162.05</b>	<b>2468</b>
		<b>161.23</b>	<b>2481</b>
0	<b>At this point, we changed the IP TOS field to 0x50 (Priority precedence, D bit set = low delay).</b>		
0	<b>LARGE</b>	<b>157.11</b>	<b>2546</b>
		<b>157.05</b>	<b>2547</b>
		<b>157.71</b>	<b>2536</b>
0	<b>We then ran without the UDP layer.</b>		
0	<b>LARGE</b>	<b>161.28</b>	<b>2480</b>
		<b>155.08</b>	<b>2579</b>
		<b>158.70</b>	<b>2520</b>

## Test Results (concluded)

- 0 The UDP layer was restored, and the IP TOS field set to 0x00 (Routine precedence, every transmission burst is ACKed at the link layer).
- 0 

<b>LARGE</b>	<b>213.88</b>	<b>1870</b>
	<b>276.42</b>	<b>1447</b>
	<b>264.50</b>	<b>1512</b>
- 0 One test was run with TOS set to 0x60 (Urgent precedence, every packet causes a coupled ACK). Transmission time was not recorded, but speed was about 960 bps for file LARGE.

# Final Observation

- 0 All 10000 bytes of MEDIUM are sent before the acknowledging GO is sent back from the receiver. Therefore, this portion of the data is being sent at the absolute maximum rate the link can achieve as configured.
- 0 The difference between the fastest 1000-byte and the fastest 10000-byte transmission was 25 seconds
  - $9000 \text{ bytes} * 8 \text{ bits per byte} / 25 = 2880 \text{ bits per second}$  with Routine precedence, D bit set for low delay
  - If we include the header bits for the 21 additional packets needed to send MEDIUM rather than SMALL, the maximum throughput is  $10260 \text{ bytes} * 8 \text{ bits per packet} / 25 \text{ seconds} = 3280 \text{ bps}$ .
- 0 This is the highest throughput that could be achieved across the link, as configured, by a "protocol" with no acknowledgements and no overhead whatsoever

# Conclusion

- 0 **The protocol stack under test achieved very close to the maximum possible throughput for a reliable transfer protocol across the network as configured.**
- 0 **The primary focus for the next phase of testing will be on performance of a modified protocol stack (with improved packet timing algorithm) under less benign network conditions:**
  - **insertion of controlled amounts of additional, conflicting, network traffic**
  - **injection of link errors**
- 0 **If possible we will operate with a more complex network, including EPLRS User Unit radios.**